arXiv:1406.4458v1 [cs.CY] 17 Jun 2014

Teaching Software Engineering through Robotics

Jiwon Shin Andrey Rusakov Bertrand Meyer
Chair of Software Engineering
Department of Computer Science
ETH Zurich, Switzerland
{jiwon.shin, andrey.rusakov, bertrand.meyer}@inf.ethz.ch

June 18, 2014

Abstract

This paper presents a newly-developed robotics programming course and reports the initial results of software engineering education in robotics context. Robotics programming, as a multidisciplinary course, puts equal emphasis on software engineering and robotics. It teaches students proper software engineering – in particular, modularity and documentation – by having them implement four core robotics algorithms for an educational robot. To evaluate the effect of software engineering education in robotics context, we analyze pre- and post-class survey data and the four assignments our students completed for the course. The analysis suggests that the students acquired an understanding of software engineering techniques and principles.

1 Introduction

Technological advancement has extended the importance of good software engineering far beyond the traditional computing devices and fields. Most university-level software engineering courses are, however, only for computer science students and do not emphasize the need for quality software in other fields. As a result, computer science students rarely gain hands-on experience with a real system, and students of other fields may only learn how to *code* but not how to *engineer* software. Robotics is one of these fields, where software engineering is an essential component, but the importance of software engineering education is overlooked.

From perception to control, robots fundamentally rely on computer programs to achieve desired behaviors. As robots advance and their complexity grows, robotics field's reliance on software libraries and thereby, its need for high-quality software augment. So far, however, roboticists have mainly focused on algorithmic development; standard software quality requirements such as modularity, correctness, and robustness have not generally been a top concern. Consequently, university-level robotics courses teach robotics algorithms but cover little to no software engineering topics.

In this paper, we present and evaluate a newly-developed, multidisciplinary course that teaches students software engineering and robotics. To our knowledge, our course, taught for the first time in the fall of 2013, is the first course to combine software engineering and robotics equally. The course teaches master's students in computer science, electrical engineering, and mechanical engineering how to bring good software engineering practices to robotics. Students apply software engineering techniques by implementing four core robotics algorithms for an educational robot. Through the course, students not only gain hands-on experience in a multidisciplinary environment but also learn to appreciate good software engineering techniques.

We evaluate the course to understand if student can learn software engineering when learning software engineering in robotics context. In particular, we investigate whether or not the students can

- learn the importance of software engineering principles, in particular, modularity and documentation

- apply software engineering techniques in programming a real system, and
- value good software engineering practice.

To answer these research questions, we conducted a pre- and post-class surveys and analyzed the four assignments students completed during the semester. The pre-class survey collected the students' educational background through multiple-choice questions. The post-survey had open-ended questions about various aspects of the course, and in this paper, we concentrate primarily on the two questions that are directly related to software engineering. The students' written responses to the questions are analyzed using thematic analysis [3]. We analyze the four assignments using software quality metrics [7] to evaluate the improvement of software quality quantitatively. Our analysis shows that students have gained an understanding and appreciation of good software engineering techniques.

The remainder of the paper is organized as follows. Section **??** presents related work. The paper continues with a description of the course in Section 2 and the methodology in Section 3. Section 4 presents the results, and Section 5 discusses them. Section 6 discusses possible threats to validity. The paper concludes with final remarks in Section 7.

The proposed robotics programming course is a multidisciplinary course that combines software engineering and robotics. The importance of multidisciplinary engineering education was noticed over a decade ago [11], and recently, Hotaling et al. have shown the benefits of multidisciplinary engineering education in preparing undergraduate students for their future [8]. In particular, as a multidisciplinary field, robotics has gained attention as a medium to teach science and engineering concepts [2]. Weinberg et al. propose to use robotics to provide hands-on instruction in various fields of engineering and computer science education and overcomes the challenges of teaching this multidisciplinary course through multidisciplinary cooperation [21]. Similarly, our course reduces the multidisciplinary difficulty by having multidisciplinary the teaching staff: of the two instructors and two teaching assistants, one instructor and one assistant have background in software engineering, and the other instructor and assistant in robotics.

Our course is a master's-level elective course that emphasizes proper software engineering and robotics algorithms. So far, however, most research on education of computer science and robotics fall into one of two categories: introductory computer science course that uses robotics as a medium to teach core computer science and/or programming concepts, or robotics courses with an emphasis on computer science that teach computational and algorithmic aspect of robotic software. In the former category of teaching introductory computer science through robotics, Fagin and Merkle show the negative effect of teaching computer science through robotics and attribute the lack of personal robots as the main reason for the failure [5]. Since then, many others have used robotics to teach introductory computer science classes with more positive outcome [19, 13, 15]. Like the successful introductory courses, our course also provides a complete robotic set to every student and allow the students to keep the hardware for the entire duration of the course. In addition, as an upper-level course, our course goes deeper into software engineering concepts such as design patterns, architecture, concurrency, and tools.

In the second category of robotics courses with computer science components, most courses focus on the computational/algorithmic aspect of robotics [9, 4]. Popular robotics textbooks may teach specific topics of robotics [20, 10] or general robotics topics with some computational principles [18], but they hardly mention of software engineering aspect of the algorithm implementation. Although Gustafson proposed using robotics to teach software engineering [6], most robotics courses teach students *what* to implement but overlook the question *how* to implement them.

To our knowledge, the only one other course, offered once in 2010 at GeorgiaTech, specifically targets to teach students how software engineering applies to robotics [1]. While the general objectives of the course is similar to ours, the course content differs significantly. GeorgiaTech course focuses mostly on software engineering with robotics as a medium to teach software engineering concepts. The course's coverage on computational aspect of robotics is limited to robot control and navigation. Our course, on the other hand, makes equal emphasis on both software engineering and robotics. We cover software engineering and robotics top-

ics equally, and the assignments require students to implement four core robotics algorithms: control and obstacle avoidance, localization, path planning, and object recognition. Students can therefore learn the core aspects of robotics while applying software engineering principles into practice.

We evaluate the course's ability to teach software engineering by combining qualitative analysis and quantitative analysis. Among various data collection techniques for field studies of software engineerings [12], we conduct questionnaires (surveys) in the beginning and at the end of the course and analyze the four assignments the students submitted during the course. The pre-survey had multiple choice questions and the post-survey had open-ended questions. To analyze the open-ended questions, we draw the inspiration from McCartney, Gokhale, and Smith [14] and apply thematic analysis [3]. Unlike the authors, we strengthen the qualitative analysis by analyzing the student-generated code using software quality metrics [7].

2 Course Description

Software engineering is a key component of modern systems, but teaching of software engineering has been limited to students of computer science. This section describes a new, multidisciplinary robotics programming course that teaches students software engineering and robotics. The course provides computer science students an opportunity to apply their software engineering skills on a real system and teaches non-computer science students, in particular, students of robotics, proper software engineering.

2.1 Objectives

The main objectives of the course are that students gain hands-on experience by programming a small robotic system with aspects of sensing, control, and planning, and gain knowledge of

- basic software engineering principles and methods,
- the most common architectures in robotics,
- coordination and synchronization methods, and
- how software engineering applies to robotics.

As a multidisciplinary, master's-level, elective course, we intend our students, who come from computer science, electrical engineering, and mechanical engineering, to learn from one another and deepen their understanding of software engineering and robotics. To address the multidisciplinary need, the course is taught by a software engineer and a roboticist and assisted by graduate students in software engineering and robotics.

2.2 Demography

The course is a master's-level elective course. The enrolment is open to any student of computer science, mechanical engineering, and electrical engineering, with some programming experience. Due to its resource intensive nature, the course is limited to 16 students. The first offering had a total of 12 students, 11 of who completed the course successfully. Of the 11 students, four were computer science students, six were mechanical engineering students, and one was a electrical engineering student. Two students were bachelor students and nine were master students.

The teaching staff comes from diverse background to reflect multidisciplinary nature of the course. In particular, the main lecturers consist of a software engineer and a roboticist and the teaching assistants are a graduate student in software engineering and a graduate student in mechanical engineering specializing. Diversity in background enables the teaching staff to better understand difficulties students would face in the course and address their needs.

2.3 Content

Lecture topics for the course ensure balanced exposure to both software engineering and robotics. Software engineering lectures include concurrency, design patterns, modern software engineering tools, and software architecture in robotics. Robotics lectures are robot control and obstacle avoidance, localization and mapping, path planning, and object recognition. In addition, there is a lecture on Robot Operating System (ROS) [17], a popular middleware in robotics, and our robotics framework.

2.4 Assignments

The course contains five assignments: an ungraded assignment for setting up the environment and four graded assignments, each implementing a core robotics algorithm. The initial ungraded assignment gives students a chance to get familiar with the hardware and software environment. The four graded assignments require students to implement algorithms for control and obstacle avoidance, localization, path planning, and object recognition. The first two graded assignments are individual assignments to encourage every student to learn the basics. The last two graded assignments are completed as a team of two students to allow the students to work on more extensive tasks.

Object-oriented programming and concurrency are the key programming paradigms for the class. As the object-oriented programming language, the course uses Eiffel and C++. The first assignment is done entirely in Eiffel, and the remaining three assignments have both Eiffel and C++ components. Eiffel with its simple concurrency extension, SCOOP, enables students to learn and program concurrent software easily [16]. C++, as a popular programming language in robotics and one of the main languages of ROS [17], enables students get familiar with a common robotics programming environment.

The course emphasizes correct implementation of robotics algorithms as well as proper software engineering. To emphasize the importance of functionality and quality, the assignments are graded equally on the demonstration of the implemented algorithm (50%) and the software quality (50%). The students submit their software via their SVN repository and demonstrate in class how their implementation works on their robot. The grading scheme for the software quality portion is as follows:

- Choice of abstraction and relations (30%)
- Correctness of implementation (40%)
- Extendibility and reusability (20%)
- Comments and documentation (10%)

To ensure continuous learning throughout the course, students receive two types of feedback after each assignment: in-class feedback and individual feedback. After each assignment, students learn in class about common mistakes and ways to avoid them in the future. In addition, we conduct individual feedback sessions. During the 15-minute long individual feedback session, students receive feedback on their particular software's correctness and quality. Immediate feedback enables students to incorporate newly-acquired knowledge into subsequent assignments and improve their software practice over time.

3 Methodology

This pilot study investigates if students can learn software engineering principles and techniques in a robotics programming class. In particular, we are interested in investigating if our students can learn the importance of software modularity and documentation. To answers these research questions, we conduct a pre-class and a post-class survey and analyze the four graded assignments the students completed for the course. The data are collected from a single-offering of the course, held in Fall 2013. This section describes the study design, data collection, and analysis method.

3.1 Surveys

We conducted pre-class survey and post-class surveys to evaluate effectiveness of teaching software engineering in the robotics programming course. The pre-class survey collected the students' educational background in software engineering, robotics/control theory, and computer vision as multiple-choice questions. The post-class survey had open-ended questions about software engineering, robotics, and overall experience. Neither survey was anonymous, and no incentive was given for completing the survey.

In this pilot study, we are particularly interested in investigating if the students understand the value of proper software engineering. We therefore analyze a subset of the survey responses that are related to software engineering. In the pre-class survey, we analyze the students' background in software engineering, i.e., programming, object-oriented programming, and concurrent programming experience and knowledge of software engineering concepts and tools. In the post-class sur-

vey, we analyze the responses to the following two questions:

- Which software engineering concepts and tools did you use in this course? Any examples?
- Has your approach to the assignment changed over the course of the semester? Do you think about your software's architecture before you begin the implementation?

We collected data from all the students who completed the course successfully. The course began with its full capacity of 16 students, but after the first two weeks of the semester, the number of students was down to 12. Of the 12 students, 11 completed the course successfully. The pre-class survey was conducted in class during the second week of the semester, and of the 11 successful students, 10 filled it out. The post-class survey was sent to the students as e-mail at the end of the semester. The students had a week to fill out and return the survey. With a friendly reminder, all 11 students returned the post-class survey.

We analyze the responses to the multiple-choice questions of the pre-class survey by counting the number of occurrences. To understand the responses to the open-ended questions of the post-class survey, we employ thematic analysis [3] to capture the main themes in the responses. We follow the procedure as suggested by Braun and Clarke and identify the themes by familiarizing ourselves with the responses, coding text fragments, grouping recurring fragments, then combining related groups into themes. To represent the underlying data correctly, we ignore rarely-occurring text fragments and verify that themes and sub-themes capture frequently occurring text fragments.

3.2 Assignments

The main objective of analyzing the assignment is to understand how the students improved quality of their software over the semester. To evaluate, we analyze the four graded assignments the students completed and submitted during the semester. Each assignment required the students to implement a core robotics algorithm, and the students received their grades based on in-class demonstration of the implemented algorithm and quality of the implemented software. In this pilot study, we investigate the change in the software quality over the four graded assignments.

The students submitted their assignments via their Subversion (SVN) repository. The first two assignments were individual assignments while the last two assignments were done in a team of two. Students had three weeks to complete the first assignment, four weeks for the second assignment, two weeks for the third assignment, and two and a half weeks for the final assignment. Varying duration reflects difference in difficulty among the assignments.

To evaluate the improvement of software quality over time, we use software metrics. Software metrics enable quantitative analysis of software. Many metrics have been suggested for measuring software quality, but no consensus exists [7]. Of many metrics, we pick a subset of the commonly-used metrics. In addition, we select some metrics to capture common mistakes of our students. Metrics used for analysis are percentage of comments, number of arguments per routine, percentage of routines with hard-coded values, lack of parametrization, and number of SVN commits. The relation between the metrics and software quality is as follows:

- Although some argue that having comments is a sign of bad software, we believe that comments are essential to reusable software. We consider software with a higher percentage of comments more reusable, but software with more lines of comments than code less reusable.
- Number of arguments per routine is also related to reusability. We consider routines with fewer arguments more reusable.
- Hard-coded values reduce reusability and extendibility. We consider software with fewer or no hard-coded values to be more reusable and extendible.
- Ability to set parameters outside of a routine or a class makes software reusable and extendible. We consider software without parametrization less reusable and extendible.
- SVN commits reflect software engineering practice, and we consider small, frequent commits to be better than large, infrequent commits.

The set of metrics we select for this analysis is neither complete nor the most representative metrics for software quality. We select these metrics for their measurability and clear link to software quality. We believe that software that perform better on these metrics is of higher quality than software that do not perform as well on them, however marginal the improvement may be.

4 Results

This section presents the results of the pre- and post-class surveys and the analysis of the assignments.

4.1 Pre-class Survey

Of the 11 students who completed the course successfully, only 10 students filled out the pre-survey. Eight were master's students, and two were upper-bachelor's students. Six students were mechanical engineering students, three were computer science students, and one was an electrical engineering student. Only one of the eight master's students had completed his bachelor's degree at ETH Zurich. The student whose data are not included in the pre-class survey results was a master's-level computer science student with a bachelor's degree from ETH Zurich.

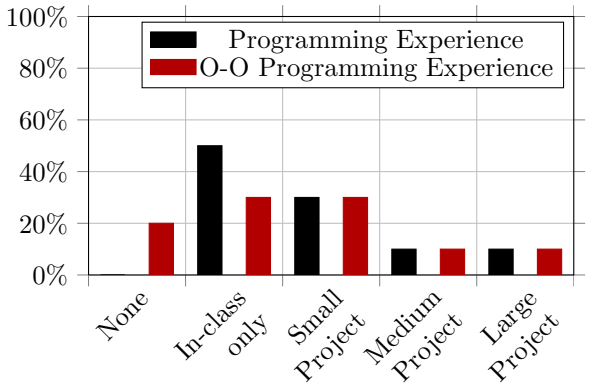

Figure 1: Prior programming experience (black) and object-oriented programming experience (red) of the students. N=10

Fig. 1 shows the distribution of programming and object-oriented programming experience our students had prior to starting the course. The course require some programming experience, and 80% of the students indicated that their programming experience is limited to class assignments or small projects consisting of less than 10'000 lines of code. In terms of prior object-oriented programming experience, 20% of the students had none, and 60% had little experience. This meant that most students had learn object-oriented paradigm before they could complete the assignments.

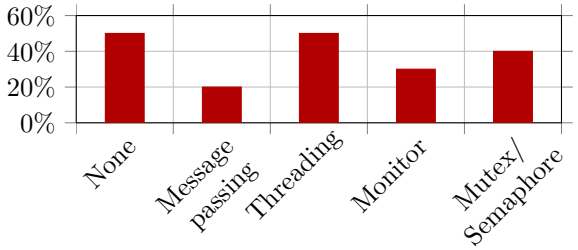

Figure 2: Experience in concurrency. N=10

In terms of concurrency (Fig. 2), half of the students had no experience while the other half had some. Of the people with some concurrent programming experience, everyone had worked with threads and most had used mutexes and/or semaphores. Every computer science student had some concurrent programming experience while most mechanical engineering student had none.

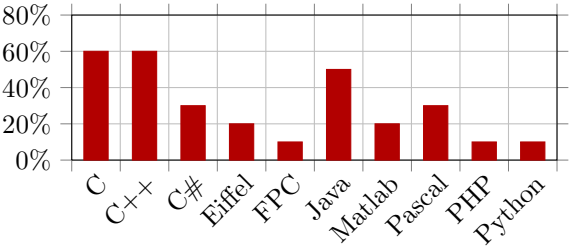

Figure 3: Familiar programming languages. N=10

The students had experience with 10 different programming languages (Fig. 3). C, C++, and Java were the most used. More than half of the students had programmed at least one of the aforementioned three languages. On average, computer science students knew 4.3 languages while other students knew 2.4 languages.

To the question on their exposure to software engineering concepts at a university, seven students indicated having learned one or more concepts (Fig. 4). One student answered not having learned any of the concepts, and two did not answer the question. On average, computer science

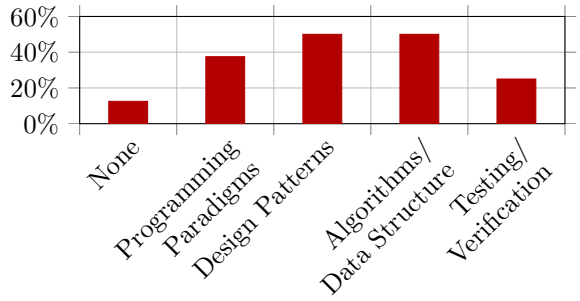

Figure 4: Software engineering concepts. N=8

students had been exposed to 3.5 concepts while other students were familiar with 1.2 concepts.

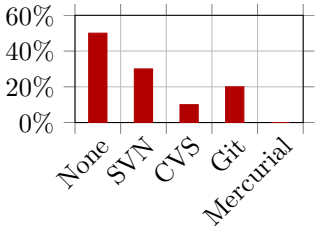

(a) Configuration management

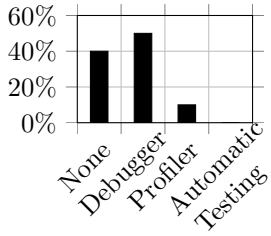

(b) Software development

Figure 5: Tools the students have used. N=10

Fig. 5 shows the students' experience with tools. Half of the students had never used any configuration management tools, and others had used SVN or Git. One student indicated that he has used both SVN and CVS. In terms of software development, four students indicated that they have not used any software development tools. Five students indicated that they have used a debugger, and one student said that he has used both a debugger and a profiler. No one had any prior experience using an automatic testing tool.

4.2 Post-class Survey

We conducted the post-class survey to understand what the students had learned in the course. All

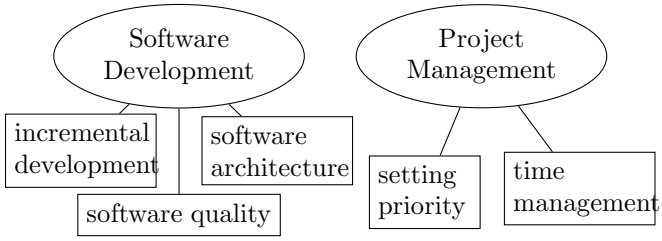

Figure 7: Change in software engineering practice with main themes (ellipse) and sub-themes (rectangle).

eleven students completed the survey. We analyze their responses to two questions: concepts and tools they used, and change in their approach to software engineering. Their written responses are grouped into themes using thematic analysis [3].

Fig. 6 shows the themes that we identified from the students' responses on concepts and tools they used. The three main themes are software quality, software tools, and programming paradigm. In software quality, the students mainly used reusability, abstraction, and documentation. In addition, code conventions, refactoring, extendibility, and readability got mentioned by one or two students. The students used editor and debugger as software tools. In programming paradigm, the students indicated using object-oriented programming and concurrency.

To the question about the change in their approach to the assignments over the course of the semester, 10 out of 11 students responded that they have changed their approach. As shown in Fig. 7, we identified two themes: software development and project management. The students improved their software development skills by thinking about software architecture, software quality, and/or applying incremental development. Many also improved project management skills in terms of time management and/or setting priority. Only one student, a master's student in computer science, indicated that there was no fundamental change to his work style as he "always plan(s) the interface first".

4.3 Assignments

We analyze the four assignments the students completed during the semester to investigate if the students learned to apply good software engineering techniques. In particular, as our course emphasized

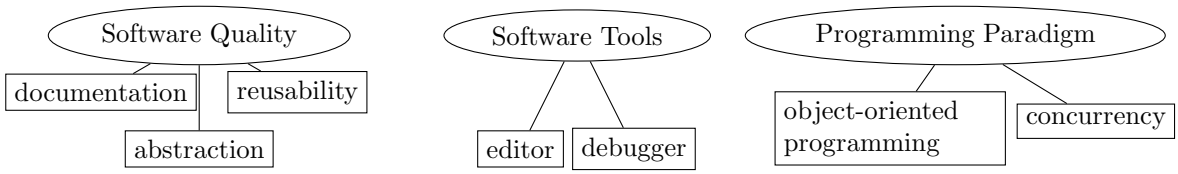

Figure 6: Concepts and tools the students used in the class. We identified three main themes (ellipse) and several sub-themes (rectangle).

software modularity and documentation, we evaluate the assignments using software quality metrics that are directly related to modularity and documentation. In addition, in grading the assignments, we identified some common mistakes that the students made:

- *Single responsibility principle* violation. We observed the tendency of having a single big class with multiple unrelated functionalities, instead of separating them into several independent classes. This results in a seriously restriction of software reusability.
- *Unnecessary dependencies* between classes. Student-generated software often contained groundless relations between classes. These unnecessary dependencies limit further extension and reuse of the software.
- Lack of *parametrization*. We observed that for the first assignment, majority of the students implemented algorithms without taking extendibility into account and did not provide a way to set parameters outside of a routine or a class. This flaw not only restricts further reusability and extendability of the system but also makes testing and debugging more complicated.
- *Hard-coded* variables. Another common mistake, which directly affects both readability and reusability, was duplication of the same values in the code. The corresponding recommendation can be formulated as avoiding so-called "hard-coded" values and "magic numbers" by introducing variables and using language support for *constants*. Fig. 11 shows how the number of this mistakes have decreased at the end of the course.

Based on the observations, we add two more metrics to our evaluation. The final set of metrics are percentage of comments, number of arguments per routine, percentage of routines with hard-coded values, lack of parametrization, and number of SVN commits.

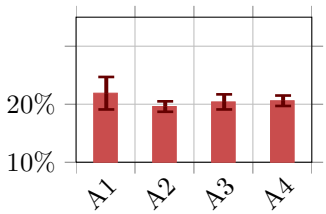

(a) Eiffel

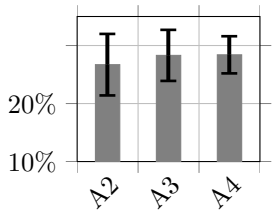

(b) C++

Figure 8: % of comments.

Comments improve understandability of code, which leads to increased reusability. As a metric for measuring comments, we compute the percentage of comments. The percentage of comments is the number of commented lines divided by the number of lines of code. This measure assumes that the code being analyzed is clean, but initially, many students submitted code without thorough clean-up. For accurate measurement, we ignored the commented out code blocks and only count the real comment lines.

Fig. 8 shows the percentage of comments for the four assignments. The students commented on average once every six to ten lines of code. The percentage of comments was higher for the first assignment than the other assignments in Eiffel, which may be due to the students reusing and modifying the example code we provided. The percentage of comments from the second assignment to the fourth assignment increased 1% in Eiffel and 1.7% in C++.

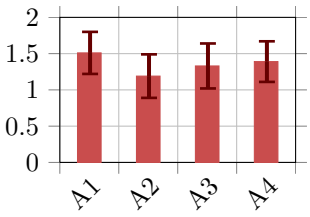

(a) Eiffel

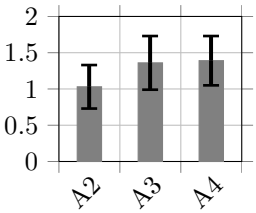

(b) C++

Figure 9: Arguments per routine.

Fig. 9 shows the average number of arguments for each assignment. Because every assignment required the students to implement a different algorithm, a direct comparison of the average number of arguments between assignments is not possible. Later assignments had algorithms with many more parameters than the earlier assignments, and an algorithm with more parameters would lead to more arguments per routine. It is important to note that despite varying number of arguments, the average number of arguments per routine remain low throughout the assignments.

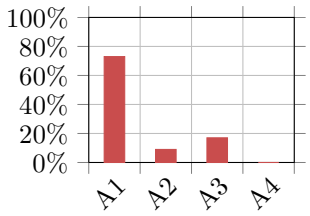

(a) Eiffel

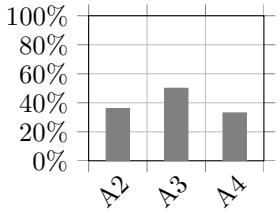

(b) C++

Figure 10: Lack of parametrization. N=11 for assignments 1 and 2. N=6 for assignments 3 and 4.

We also counted the number of submitted solutions that lacked parametrization. The results are shown in Fig. 10. In the first assignment, over 70% of the solutions lacked parametrization. After we pointed out the mistake during the feedback phase, the occurrence dropped significantly.

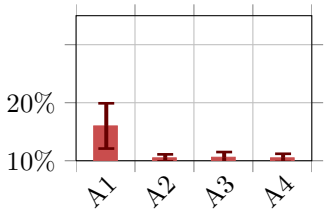

(a) Eiffel

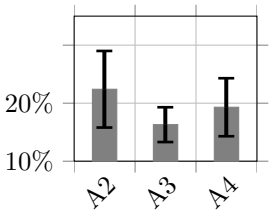

(b) C++

Figure 11: % of routines with hard-coded values.

Hard-coded values and magic numbers were also prominent in the first assignment. During the feedback session, we recommended the students to introduce variables and use language support for *constants*. Fig. 11 shows how the percentage of routines with hard-coded values changed over the assignments. There is a significant drop from the first to the second assignment in Eiffel.

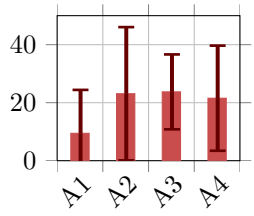

(a) Total commits

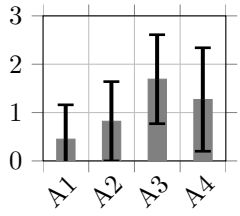

(b) Daily commits

Figure 12: SVN repository usage.

Fig. 12 shows the SVN repository usage. As half of the students had never used any configuration management tool, the initial SVN usage was relatively low. The students who were familiar with SVN committed their incremental changes while those who were new made commits sporadically or not at all. After the first assignment, however, the students started using SVN more regularly. The

number of total commits per assignment jumped after the first assignment and remained steady. The daily usage of SVN, the number of commits divided by the number of days, increases over the first three assignment before dropping for the fourth assignment. The drop may have been caused by time crunch at the end of the semester.

5 Discussion

This section discusses meaning of the results in terms of the students' knowledge gain in software engineering.

5.1 Pre-class survey

The pre-survey revealed that majority of the students came to the class with limited experience in programming and little to no exposure to key software engineering tools. The discrepancy in software engineering knowledge and experience between computer science students and non-computer science students was quite stark. Despite the importance of software engineering in other fields, some non-computer science students had only programmed in class and had no exposure to object-oriented programming. Given that the course is multidisciplinary and hence requires knowledge in many different fields, educational background and programming experience of a student had little effect on the student's final performance in the class.

5.2 Post-class survey

We analyzed two open-ended questions of the post-survey. Responses to the first question revealed that the students used several concepts and tools during the semester. Most students mentioned concepts related to software quality, including reusability, abstraction, and documentation. This shows that our students developed some understanding of software quality, one of our main objectives. Software tools with editor and debugger as sub-themes was also frequent. As the question contained the word *tools*, identifying tools as a theme is not surprising; however, given that we devoted a lecture on software engineering tools, we expected our students to have used more than editor and debugger. The third theme, programming paradigm, with its two sub-themes, object-oriented programming and concurrency, show that the students became more aware of these paradigms compared to the beginning of the course.

Our analysis of the second question revealed that most students improved software development and/or project management. In software development, the students indicated that they now think about software architecture and quality and develop software incrementally. This result was as we had hoped. What we did not expect as much was their improvement in project management skills. Many indicated that they now manage their time better and set priority when working. This is a positive by-product of the students working on challenging assignments with limited time, and these skills are definitely useful in software engineering.

In general, we faced several issues in analyzing the responses. One major issue is that some students gave long, detailed answers to the survey questions whereas others did not. Short responses contain, inevitably, limited information. While it is likely that the items mentioned in the response are those that the students found most important, it is difficult to judge if the response captures the complete set. Other issue is that the questions were narrowly formulated. While this reduces potential misunderstanding, it also limits variety of the responses. This was especially true for the second question, to which most students gave yes or no answer with a short explanation. Last, the post-survey contained many other questions not reported in this paper, and the sheer number of question may have caused some students to provide only a brief response to each question.

5.3 Assignments

Quantitative analysis of the assignments using software quality metrics revealed that the students improved after the first assignment in some areas and remained the same in others. Among the five metrics we used for the analysis, percentage of comments and number of arguments per routine did not reveal much in terms of improvement of the students' software engineering techniques. Parametrization, hard-code variables, and SVN usage showed that the students made a leap after the first assignment. Likely cause of improvement is the individual feedback session. After the first feedback

session, many students remarked that no one had given them the kind of tips we gave. In the subsequent assignments, quality of student-generated software depended more on difficulty of the assignments and students' willingness to put their knowledge into practice. When we pointed out mistakes during the assignment 2 and 3 feedback sessions, the students often showed their awareness of the mistakes and blamed the time-constraint for not applying their knowledge into practice.

The analysis revealed that in general, the code written in Eiffel were of higher quality than those in C++. This may be because for each assignment, C++ portion required the students to implement a new robotics algorithm whereas Eiffel portion required them to extend the previous version. We also observed high variability among the submitted software. Some students submitted meticulously cleaned-up, refactored, and commented code while others submitted code with commented-out code blocks and To-Do notes. Despite the variability, most students improved after the first feedback session.

5.4 General remark

Relating the survey results and the analysis of the assignments shows that our students came in with limited knowledge of software engineering and left with an understanding of software engineering principles and techniques. The course focused on modularity and documentation as the main indicators of quality software, and the survey responses revealed that the students' appreciation of software quality is directly related to these. Minimization of hard-coded variables and introduction of parametrization in their assignment reflect that the students tried to put their knowledge into practice. The students also claimed that they had changed to incremental development, and increase in the frequency of SVN commits is in line with their statement. Overall, the students may not have internalized all the pitfalls, but they have gained an understanding of proper software engineering and applied some into practice.

6 Threats to Validity

There are several limitations and threats to validity of this study. An obvious limitation of the study is that data are drawn from a single course offering at one university. While the pilot study provides some understanding of the effect of teaching software engineering in a robotics programming course, longer and broader study is necessary to generalize the claim. In addition, the study's small data size limits generalization of the results.

The study contains several potential sources of bias. First, as an elective course, no student was required to take the course. The students may have been more motivated to learn about software engineering in robotics setting than regular students, and this may bias our results towards better learning outcome. Second, the authors designed, implemented, and executed the course described in the paper. Even though we tried to be neutral in analyzing data, the results may nonetheless be biased towards success. Lastly, as the students got to know the authors very well by the end of the course, their post-survey responses may have been influenced by their feeling towards the authors. In fact, top third of the students gave longer and more-detailed responses to the survey than most other students, resulting in their opinion being more reflected in the qualitative analysis. This may bias towards better post-survey results.

In our analysis, we assume that the assignments capture their understanding of software engineering. In reality, the correlation between the submitted code and students' knowledge is not clear. Many students found the course challenging and time-consuming and tried to optimize the way they spend their time. Consequently, software they submitted is not necessarily representative of what they actually learned in class. Due to time limitation, students may have not applied everything they have learned but instead only those that result in better grade for the effort. In addition, the last two assignments were completed in a team of two students, and it is not clear how the knowledge of two students maps to a single piece of software.

7 Conclusions

This paper reported and evaluated a newly-developed robotics programming course that emphasizes proper software engineering – modularity and documentation – in robotics context. The course covers topics of software engineering and robotics equally, and the students learn to apply proper software engineering techniques by implementing four different robotics algorithms for an educational robot. To understand the effect of the course, we conducted a pre-survey and a post-survey and analyzed the four assignments using software quality metrics. The analysis of survey responses and student-generated code showed that the students gained some understanding of software engineering principles and techniques.

Based on our analysis in this pilot study, we believe that students *can* learn software engineering by programming a robot, and software engineering education is important to students outside of computer science. Fully understanding the effect of teaching software engineering in a robotics setting requires more extensive research. To validate the results of this pilot study, we plan to continue our evaluation of the course in the subsequent offerings and collect data over a longer period. It would be especially interesting to understand the long-term effect of our course to computer science students and other students alike.

8 Acknowledgments

The research leading to these results has received funding from the European Re- search Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Grant agreement no. 291389, Hasler Foundation under SmartWorld initiative / Roboscoop: concurrent robotics framework project, and ETH Zurich under Innovedum / project no. 733.

References

[1] Software engineering in robotics, 2014. http://www.cc.gatech.edu/~hic/8803-SER-10 [Online].

[2] R. D. Beer, H. J. Chiel, and R. F. Drushel. Using autonomous robotics to teach science and engineering. *Commun. ACM*, 42(6):85–92, June 1999.

[3] V. Braun and V. Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3:77–101, 2006.

[4] N. Correll, R. Wing, and D. Coleman. A one-year introductory robotics curriculum for computer science upperclassmen. *IEEE Transactions on Education*, 56(1):54–60, February 2013.

[5] B. Fagin and L. Merkle. Measuring the effectiveness of robots in teaching computer science. *SIGCSE Bull.*, 35(1):307–311, Jan. 2003.

[6] D. Gustafson. Using robotics to teach software engineering. In *Frontiers in Education Conference*, volume 2, pages 551–553 vol.2, Nov 1998.

[7] B. Henderson-Sellers. *Object-oriented Metrics: Measures of Complexity*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1996.

[8] N. Hotaling, B. Fasse, L. F. Bost, C. D. Hermann, and C. R. Forest. A quantitative analysis of the effects of a multidisciplinary engineering capstone design course. *Journal of Engineering Education*, 101(4):630–656, October 2012.

[9] J.-F. Lalonde, C. Hartley, and I. Nourbakhsh. Mobile robot programming in education. In *Proceedings of ICRA '06*, pages 345–350, May 2006.

[10] S. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, UK, 2006.

[11] E. A. Lee and D. G. Messerschmitt. Engineering an education for the future. *Computer*, 31(1):77–85, Jan. 1998.

[12] T. C. Lethbridge, S. E. Sim, and J. Singer. Studying software engineers: Data collection techniques for software field studies. *Empirical software engineering*, 10(3):311–341, 2005.

[13] S. A. Markham and K. N. King. Using personal robots in cs1: Experiences, outcomes, and attitudinal influences. In *ITiCSE '10*, pages 204–208, 2010.

[14] R. McCartney, S. S. Gokhale, and T. M. Smith. Evaluating an early software engineering course with projects and tools from open source software. In *Proceedings of ICER '12*, pages 5–10, 2012.

[15] M. M. McGill. Learning to program with personal robots: Influences on student motivation. *Trans. Comput. Educ.*, 12(1):4:1–4:32, Mar. 2012.

[16] S. Nanz, F. Torshizi, M. Pedroni, and B. Meyer. Design of an empirical study for comparing the usability of concurrent programming languages. In *Proceedings of ESEM'11*, pages 325–334. IEEE Computer Society, 2011.

[17] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng. Ros: an open-source robot operating system. In *IROS '09*, 2009.

[18] R. Siegwart, I. R. Nourbakhsh, and D. Scaramuzza. *Introduction to Autonomous Mobile Robots*. MIT Press, Cambridge, MA, USA, 2nd edition, 2011.

[19] J. Summet, D. Kumar, K. O'Hara, D. Walker, L. Ni, D. Blank, and T. Balch. Personalizing cs1 with robots. *SIGCSE Bull.*, 41(1):433–437, Mar. 2009.

[20] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, Cambridge, MA, USA, 2005.

[21] J. B. Weinberg, G. L. Engel, K. Gu, and C. S. Karacal. A multidisciplinary model for using robotics in engineering education. In *Proceedings of the 2001 American Society for Engineering Education Annual Conference and Exposition*, 2001.